On Affine Invariant L_p Depth Classifiers based on an Adaptive Choice of p

Subhajit Dutta[1] and Anil K. Ghosh[2]

November 18, 2016

arXiv:1611.05668v1 [stat.ME] 17 Nov 2016

Abstract

In this article, we use L_p depth for classification of multivariate data, where the value of p is chosen adaptively using observations from the training sample. While many depth based classifiers are constructed assuming elliptic symmetry of the underlying distributions, our proposed L_p depth classifiers cater to a larger class of distributions. We establish Bayes risk consistency of these proposed classifiers under appropriate regularity conditions. Several simulated and benchmark data sets are analyzed to compare their finite sample performance with some existing parametric and nonparametric classifiers including those based on other notions of data depth.

Key words: Bayes risk, Data depth, Kernel density estimation, l_p-symmetric distributions, Maximum likelihood estimation, Misclassification rates.

Short title: Affine invariant L_p depth classifiers

[1]Department of Mathematics and Statistics, Indian Institute of Technology, Kanpur - 208016, India.
E-mail: duttas@iitk.ac.in

[2]Theoretical Statistics and Mathematics Unit, Indian Statistical Institute, 203, B. T. Road, Kolkata - 700108, India.
E-mail: akghosh@isical.ac.in

1 Introduction

Data depth measures the centrality of a point $\mathbf{x}$ in $\mathbb{R}^d$ with respect to a d-dimensional data cloud, or a d-dimensional probability distribution. As a result, it provides a centre-outward ordering of multivariate data. Various notions of data depth are available in the literature (see, e.g., Liu, Parelius and Singh, 1999; Zuo and Serfling, 2000), and they have been used for generalizing many univariate statistical methods to the multivariate setup. One important application is supervised classification. Ghosh and Chaudhuri (2005) introduced maximum depth classifiers and also developed a modified classifier based on Tukey's (1975) half-space depth (HD). Later, Dutta and Ghosh (2012) investigated some robust classifiers based on projection depth (PD) (see, e.g., Zuo and Serfling, 2000) and robust versions of Mahalanobis depth (MD). Li, Cuesta-Albertos and Liu (2012) developed nonparametric classifiers based on depth-depth (DD) plots. Other depth based classification methods include the work of Jörnsten (2004), Hoberg and Mosler (2006), Hartikainen and Oja (2006), Cui, Lin and Yang (2008) and Paindaveine and Van Bever (2015).

Constructions of most of these classifiers were motivated by elliptic symmetry (i.e., l_2-symmetry after an affine transformation; see, e.g., Fang, Kotz and Ng, 1989) of the competing class distributions. In this article, we develop some classifiers motivated by general l_p-symmetry (after an affine transformation) of the underlying distributions. We assume that the density of each competing class to be continuous, and of the form $f(\mathbf{x}) = \psi(\|\mathbf{A}(\mathbf{x}-\mathbf{b})\|_p)$, where $p \geq 1$, $\mathbf{b}$ is a d-dimensional vector, $\mathbf{A}$ is a $d \times d$ non-singular matrix and $\psi : \mathbb{R}_+ \to \mathbb{R}_+$ is a scalar continuous function. Here, $\|\mathbf{z}\|_p = (|z_1|^p + \cdots + |z_d|^p)^{1/p}$ for any $\mathbf{z} = (z_1, \ldots, z_d)^T \in \mathbb{R}^d$. Clearly, the location (i.e., the centre of symmetry) $\boldsymbol{\mu}$ of this distribution is $\mathbf{b}$. In the case of $p = 2$ (i.e., elliptic symmetry), the associated scatter matrix $\boldsymbol{\Sigma}$ is given by $(\mathbf{A}\mathbf{A}^T)^{-1}$. If f is assumed to have finite second moments, it has the mean vector $E[\mathbf{X}] = \mathbf{b}$ and the dispersion matrix $E[(\mathbf{X}-\boldsymbol{\mu})(\mathbf{X}-\boldsymbol{\mu})^T] = c_{p,\psi}\,(\mathbf{A}\mathbf{A}^T)^{-1}$, where $c_{p,\psi}$ is a positive constant that depends on p and ψ. So, $\mathbf{A}$ can be viewed as a square root of $\boldsymbol{\Sigma}^{-1}$ (upto a scalar constant). Several authors have studied various properties of multivariate l_p-symmetric distributions (see, e.g., Yue and Ma, 1995; Gupta and Song, 1997). Sinz and Bethge (2010) used l_p-symmetric distributions to develop various statistical tools including independent

component analysis. Arellano-Valle and Richter (2012) have introduced skewed versions of l_p-symmetric distributions. Throughout this article, by l_p-symmetric, we mean l_p-symmetric after applying an affine transformation (as we have discussed above).

If the underlying class distributions are l_2-symmetric (i.e., spherically or elliptically symmetric), for several existing notions of data depth, the class densities turn out to be functions of depths. Consequently, the Bayes classifier can be expressed as a function of data depths corresponding to different competing classes. This is the main argument used in proving the Bayes risk consistency for most of the existing depth based classifiers. However, if the depth contours fail to match the density contours, one cannot use such mathematical arguments. To appreciate this, consider a two-class problem with equal priors and class densities $f_1(\mathbf{x}) = c_p \exp(-\|\mathbf{x}\|_p^p)$ and $f_2(\mathbf{x}) = c_p/\sigma^d \exp(-\|\mathbf{x}\|_p^p/\sigma)$, where $\sigma \neq 1$ is a positive constant, and c_p is the normalizing constant. Clearly, the boundary of the Bayes classifier is of the form $\{\mathbf{x} : \|\mathbf{x}\|_p = K_0\}$ for some $K_0 > 0$. Now, MD contours for each class are concentric hyper-spheres of the form $\{\mathbf{x} : \|\mathbf{x}\|_2 = K\}$ for varying choices of $K > 0$. Therefore, when $p \neq 2$, no classifier based on MD can yield the Bayes class boundary.

For l_p-symmetric distributions with $p \neq 2$, Dutta, Ghosh and Chaudhuri (2011) proved that the density cannot be a function of HD as well. Figure 1 shows HD and PD contours (indicated using `bold curves`) computed based on 2000 observations from two l_p-symmetric densities (density contours are shown using `dotted curves`) with $p = 1$ and $p = 5$ (see the left and the middle panels). This figure clearly shows that the class density cannot be a function of depth in either of these cases.

To overcome this limitation, we use L_p depth (see, e.g., Zuo and Serfling, 2000) with a data driven choice of p. L_p depth ($\mathrm{L}_p\mathrm{D}$) of an observation $\mathbf{x}$ with respect to a multivariate distribution function F is defined as $\delta_p(\mathbf{x}) = [1 + r_p(\mathbf{x})]^{-1}$, where $r_p(\mathbf{x}) = \|\mathbf{\Sigma}^{-1/2}(\mathbf{x} - \boldsymbol{\mu})\|_p$ with $\boldsymbol{\mu}$ and $\mathbf{\Sigma}$ being the location and the scatter associated with F. The empirical version of $\delta_p(\mathbf{x})$ is obtained by estimating p, and $r_p(\mathbf{x})$ from the data (see Section 2 for details). The right panel of Figure 1 shows that the empirical $\mathrm{L}_p\mathrm{D}$ contours based on 2000 observations almost coincide with the underlying l_p-

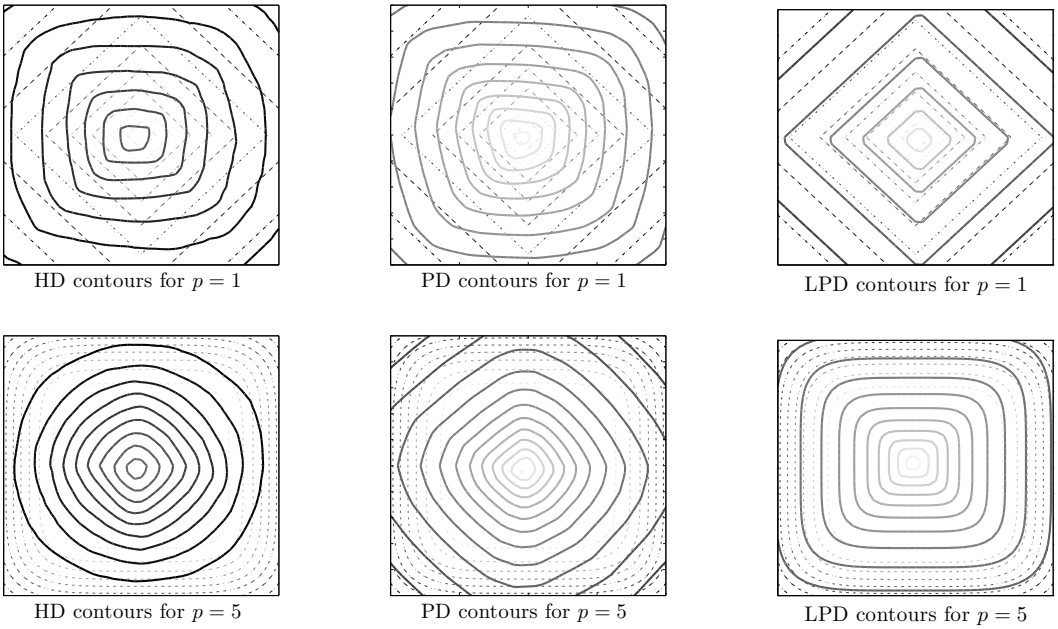

Figure 1: Density contours (`dotted curves`) of l_p symmetric distributions for $p = 1$ and $p = 5$, and the corresponding HD, PD and L_pD contours (`bold curves`) estimated from the data.

symmetric density contours both for $p = 1$ and $p = 5$. So, classifiers based on $\delta_p(\mathbf{x})$ are expected to yield an improved performance.

2 Estimation of L_p Depth

To construct a classifier based on $\delta_p(\mathbf{x})$, first one needs to find an appropriate value of p for each of the competing classes. We choose the value of p which fits to the data well. Suppose $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n$ are n independent observations from a l_p-symmetric distribution with density f. If $\mathbf{b}$ and $\mathbf{A}$ denote the associated location and scatter parameters (defined in Section 1), then f can be expressed as:

$$f(\mathbf{x}, p) = |\mathbf{A}| \, \frac{p^{d-1}\Gamma(d/p)}{2^d[\Gamma(1/p)]^d} \, \frac{g_p(\delta_p(\mathbf{x}))\, \delta_p(\mathbf{x})^{d+1}}{[1-\delta_p(\mathbf{x})]^{d-1}},$$

where $\delta_p(\mathbf{x}) = [1 + \|\mathbf{A}(\mathbf{x} - \mathbf{b})\|_p]^{-1}$, and g_p is the density of $\delta_p(\mathbf{X})$ (see Lemma 1 in the Appendix). In Section 2.1, we discuss the estimation procedure for p, assuming that $\hat{\mathbf{b}}$ and $\hat{\mathbf{A}}$ (estimates for $\mathbf{b}$ and $\mathbf{A}$, respectively) are given. In Section 2.2, we construct $\hat{\mathbf{b}}$ and $\hat{\mathbf{A}}$ which makes the estimate of p, and hence the empirical version of L_pD affine invariant.

2.1 Estimation of p

For any fixed p, using the data $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n$ from f, we first compute $\hat{r_p}(\mathbf{x}_i) = \|\hat{\mathbf{A}}(\mathbf{x}_i - \hat{\mathbf{b}})\|_p$ and $\hat{\delta}_p(\mathbf{x}_i) = [1 + \hat{r_p}(\mathbf{x}_i)]^{-1}$ for $i = 1, 2, \ldots, n$. In order to estimate f, one needs an estimate of g_p. Assuming $\hat{\delta_p}(\mathbf{x}_1), \ldots, \hat{\delta_p}(\mathbf{x}_n)$ as sample observations, we estimate it using the kernel density estimation (see, e.g., Silverman, 1998) method. This density estimate is given by $\hat{g}_{p,h}(\delta) = \frac{1}{nh} \sum_{i=1}^{n} K\left[h^{-1}(\delta - \hat{\delta_p}(\mathbf{x}_i))\right]$, where K is the kernel function and h is the associated bandwidth parameter. We use the Gaussian kernel $K(t) = (2\pi)^{-1/2} e^{-t^2/2}$, and h is chosen using the bandwidth selection method proposed in Sheather and Jones (1991). So, under the assumption of l_p symmetry, the estimate of the density function f is given by

$$\hat{f}_h(\mathbf{x}, p) = |\hat{\mathbf{A}}| \, \frac{p^{d-1}\Gamma(d/p)}{2^d\{\Gamma(1/p)\}^d} \frac{\hat{g}_{p,h}(\hat{\delta_p}(\mathbf{x}))\, \hat{\delta_p}(\mathbf{x})^{d+1}}{[1 - \hat{\delta}_p(\mathbf{x})]^{d-1}}.$$

Irrespective of the dimension of the data, here we need only one-dimensional kernel density estimation. Similar approaches for depth based density estimation were also used by Fraiman, Liu and Mechole (1997) and Hartikainen and Oja (2006).

We compute the estimated joint likelihood $\mathcal{L}_p(\mathbf{x}_1, \ldots, \mathbf{x}_n) = \prod_{i=1}^{n} \hat{f}_h(\mathbf{x}_i, p)$ for different p, and choose the one that maximizes $\mathcal{L}_p(\mathbf{x}_1, \ldots, \mathbf{x}_n)$ or $\log \mathcal{L}_p(\mathbf{x}_1, \ldots, \mathbf{x}_n)$. However, note that if $\hat{\delta}_p(\mathbf{x})$ is close to zero or unity, $|\log \hat{f}_h(\mathbf{x}, p)|$ tends to be very influential. So, we consider only those $\mathbf{x}_i$s for which $\hat{\delta}_p(\mathbf{x}_i)$ lie between ζ_{1n} and ζ_{2n} $(0 < \zeta_{1n} < \zeta_{2n} < 1)$ and find $\hat{p}$ (the estimate of p) by maximizing

$$\varphi_p(\mathbf{x}_1, \cdots, \mathbf{x}_n) = \sum_{\{i \,:\, \hat{\delta}_p(\mathbf{X}_i) \in [\zeta_{1n}, \zeta_{2n}]\}} \log \hat{f}_h(\mathbf{x}_i, p)$$

over $\mathcal{P} = \{p_1, p_2, \ldots, p_M\}$, a finite set of values for p. The following theorem suggests a suitable choice for $[\zeta_{1n}, \zeta_{2n}]$, and gives the asymptotic behavior of $\hat{p}$ under appropriate regularity conditions.

Theorem 1: *Let $\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_n$ be independent and identically distributed with the density of the form $f(\mathbf{x}, p_0) = \psi_0(\|\mathbf{A}(\mathbf{x} - \mathbf{b})\|_{p_0})$ for some $p_0 \geq 1$. For any $p \geq 1$, define $\zeta_{1n} = \zeta_1(p)$ and $\zeta_{2n} = \zeta_2(p)$ as the α_{1n}-th and α_{2n}-th quantile of g_p, where $\alpha_{1n} \downarrow 0$, $\alpha_{2n} \uparrow 1$ and $n^{1/2} \min\{\zeta_{1n}, 1 - \zeta_{2n}\} \to \infty$ as $n \to \infty$. Now, consider the following assumptions:*

(C1) $\hat{\mathbf{A}}$ *and* $\hat{\mathbf{b}}$ *are* $\sqrt{n}$*-consistent estimates of* $a_0\mathbf{A}$ *and* $\mathbf{b}$ *(i.e.,* $\sqrt{n}\|\hat{\mathbf{A}} - a_0 A\|_F = O_P(1)$*, where* $a_0 > 0$ *and* $\|\cdot\|_F$ *denotes the Frobenious norm; and* $\sqrt{n}\|\hat{\mathbf{b}} - \mathbf{b}\|_p = O_P(1)$ *for any* $p \geq 1$*), respectively.*
(C2) For any $p \geq 1$*,* g_p *is absolutely continuous and the bandwidth* h *associated with the estimation of* g_p *is of the order* $O(n^{-1/4+\varepsilon})$ *for some* $\varepsilon \in (0, 1/4)$*.*
(C3) For any $p \geq 1$*,* $n^{1/2}h^2 \inf_{\{\delta_p(\mathbf{X}) \in [\zeta_{1n}, \zeta_{2n}]\}} g_p(\delta_p(\mathbf{x})) \to \infty$ *as* $n \to \infty$*.*
Assume (C1)-(C3), and define $\hat{p}_n = \arg\max_p \varphi_p(\mathbf{X}_1, \cdots, \mathbf{X}_n)$*. If* $p_0 \in \mathcal{P}$*, then* $\hat{p}_n \xrightarrow{P} p_0$ *as* $n \to \infty$*. If* $p_0 \notin \mathcal{P}$*, then* $\hat{p}_n \xrightarrow{P} p_0^*$*, where* p_0^* *minimizes the Kullback-Leibler divergence between* $f(\cdot, p_0)$ *and* $f(\cdot, p)$ *over* $p \in \mathcal{P}$*.*

If g_p is bounded away from zero and infinity, it is easy to see that $(C3)$ holds for any choice of α_{1n} and α_{2n}. In fact, $(C3)$ holds if g_p remains bounded away from zero on any bounded interval inside its support. Also note that the Sheather-Jones bandwidth (see Sheather and Jones, 1991) that we use for kernel density estimation is of the order $O_P(n^{-1/5})$.

The quantities $\hat{\mathbf{b}}$ and $\hat{\mathbf{A}}$ can be obtained directly from $\hat{\boldsymbol{\mu}}$ and $\hat{\boldsymbol{\Sigma}}$, the estimates of $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$. Affine equivariant estimates of $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ (e.g., usual moment based estimates or minimum covariance determinant (MCD) estimates) can be used for this purpose. Assuming the existence of second order moments, the sample mean converges to $\boldsymbol{\mu}$ and the sample dispersion matrix converges to constant multiple of $\boldsymbol{\Sigma}$ at $\sqrt{n}$ rate as stated in $(C1)$. Under appropriate regularity conditions, we have $\sqrt{n}$ convergence for MCD estimates as well (see Cator and Lopuhaä, 2012).

2.2 Estimation of b and A

We can continue to use $\hat{\boldsymbol{\mu}}$ (the moment based estimate, or the MCD estimate) as an estimate of $\mathbf{b}$. However, to construct $\hat{\mathbf{A}}$, one has to compute the square root of $\hat{\boldsymbol{\Sigma}}$. Note that the symmetric square root of $\hat{\boldsymbol{\Sigma}}$ is not affine eqivariant. We construct an affine equivariant square root of $\hat{\boldsymbol{\Sigma}}$ using the transformation re-transformation technique (see, e.g., Chakraborty and Chaudhuri, 1996). Consider a subset $\alpha = \{i_1, i_2, \ldots, i_{d+1}\}$ of the set $\{1, 2, \ldots, n\}$. Use it to construct a $d \times d$ matrix $\mathbb{X}(\alpha) = [(\mathbf{x}_{i_1} - \mathbf{x}_{i_{d+1}}) (\mathbf{x}_{i_2} - \mathbf{x}_{i_{d+1}}) \ \ldots \ (\mathbf{x}_{i_d} - \mathbf{x}_{i_{d+1}})]$, and compute $\mathbb{Z}(\alpha) = \mathbb{X}(\alpha)^T \hat{\boldsymbol{\Sigma}}^{-1} \mathbb{X}(\alpha)$. Now, find an α_0 that makes $\mathbb{Z}(\alpha)$ close to a constant multiple of $\mathbf{I}_d$ (the $d \times d$ identity matrix). For practical

purposes, one may choose an α_0 that maximizes the ratio of the determinant of $\mathbb{Z}(\alpha)$ to the trace of $\mathbb{Z}(\alpha)$. However, finding the actual maximizer α_0 is computationally difficult. So, following Chakraborty and Chaudhuri (1996), we choose an α which makes the ratio is close to 1 (≥ 0.99).

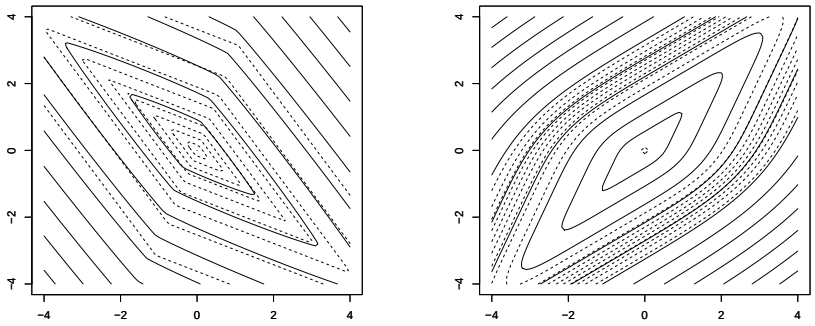

Figure 2: Density contours (`dotted curves`) and estimated L_pD contours (`bold curves`).

The matrix $\mathbb{X}(\alpha_0)$ can now be considered as a square root of $\hat{\mathbf{\Sigma}}$ (upto a scalar multiple). The use of $\hat{\mathbf{A}} = \mathbb{X}^{-1}(\alpha_0)$ makes $\hat{\delta}_p(\mathbf{x})$ affine invariant. To study the empirical performance of this estimation method, we generated 400 observations from the density $f_1(x_1, x_2) \propto \exp\{-(|x_1| + |x_2|/0.3)\}$ (which is l_1-symmetric), and then rotated them by an angle of $3\pi/4$. The left panel in Figure 2 shows that the estimated L_pD contours (solid curves) closely matches the corresponding density contours (dotted curves). We carried out a second experiment with observations from the density $f_2(x_1, x_2) \propto \exp\{-(|x_1|^5 + |x_2|^5/0.3)\}$ (which is l_5-symmetric), which were further rotated by an angle of $\pi/4$. In this example also, the estimation method worked quite well (see the right panel in Figure 2).

3 Classification with L_p Depth

In this section, we first study the maximum L_pD classifier. This classifier works well when the competing classes have the same prior, and they differ only in their locations. However, if the classes have unequal priors and/or they differ only in their scatters and shapes, the maximum depth classifier may fail to have satisfactory performance. To cope with such cases, we develop a modified classifier in Section 3.2.

3.1 Maximum L_p depth classifier

The maximum L_pD classifier classifies an observation to the class with respect to which it has the maximum depth. To construct the empirical version of L_pD, we estimate p and $\delta_p(\mathbf{x})$ from the data (as discussed in Section 2). Suppose there are J competing classes with densities $f_1, \ldots, f_J$, and for the j-th class, these estimates are denoted by $\hat{p}_j$ and $\hat{\delta}_{\hat{p}_j,j}(\mathbf{x})$, respectively, for $1 \leq j \leq J$. Then, the maximum L_pD classifier classifies an observation $\mathbf{x}$ to the j-th class if $\hat{\delta}_{\hat{p}_j,j}(\mathbf{x}) > \hat{\delta}_{\hat{p}_i,i}(\mathbf{x})$ for all $i \neq j$. However, this classifier is mainly used when the competing classes differ only in their locations, and in such cases it is more appealing to use a common value of p for all classes. If $\mathbf{x}_{j1}, \mathbf{x}_{j2}, \ldots, \mathbf{x}_{jn_j}$ are observations from the j-th ($1 \leq j \leq J$) class, this common value of p can be estimated by maximizing the joint log-likelihood function $\sum_{j=1}^{J} \varphi_p(\mathbf{x}_{j1}, \mathbf{x}_{j2}, \ldots, \mathbf{x}_{jn_j})$ over $\mathcal{P}$. The resulting maximum L_pD classifier is given by $\mathsf{d}_1(\mathbf{x}) = \arg\min_{1 \leq j \leq J} \hat{\delta}_{\hat{p},j}(\mathbf{x})$, where $\hat{p}$ denotes the common estimated value of p.

Theorem 2: *Assume the density functions $f_1, \ldots, f_J$ to be unimodal, and f_j is of the form $f_j(\mathbf{x}, p_0) = \psi(\|\mathbf{A}(\mathbf{x} - \mathbf{b}_j)\|_{p_0})$ for some $p_0 \in \mathcal{P}$ ($1 \leq j \leq J$). If the prior probabilities of the competing classes are equal, under conditions* ($C1$) - ($C3$) *(stated in Theorem 1), the misclassification rate of* $\mathsf{d}_1(\mathbf{x})$ *converges to the Bayes risk as $min\{n_1, \ldots, n_J\} \to \infty$.*

3.2 Generalized L_p depth classifier

In practice, the prior probabilities of different classes may not be equal, and the class distributions can also differ in their scatters and shapes. In such cases, maximum depth classifiers may have higher misclassification probabilities (see, e.g., Dutta and Ghosh, 2012). We now construct the generalized L_pD classifier to cope with such situations. Suppose that there are J competing classes, and $f_{h_j}(\mathbf{x}, \hat{p}_j)$ is the density estimate for the j-th class ($1 \leq j \leq J$). In a two-class problem, this generalized L_pD classifier is given by

$$\mathsf{d}_2(\mathbf{x}) = \begin{cases} 1 & \text{if } \hat{f}_{h_1}(\mathbf{x}, \hat{p}_1) / \hat{f}_{h_2}(\mathbf{x}, \hat{p}_2) > k, \\ 2 & \text{otherwise.} \end{cases}$$

where k is chosen by minimizing the leave-one-out cross-validation estimate of the misclassification probability. For more than two classes, we use the pairwise classification approach followed by the method of majority voting.

Theorem 3: *Suppose that for all* $1 \leq j \leq J$*, the density* f_j *is continuous, it has support over entire* $\mathbb{R}^d$ *and is of the form* $f(\mathbf{x}, p_j^\circ) = \psi_j(\|\mathbf{A}_j(\mathbf{x} - \mathbf{b}_j)\|_{p_j^\circ})$ *for some* $p_j^\circ \in \mathcal{P}$*. Also, assume that the optimal Bayes classifier has non-empty favorable regions for each of the* J *classes. Then, under the conditions* $(C1)$ - $(C3)$ *(stated in Theorem 1), the misclassification rate of the generalized* L_pD *classifier* $\mathsf{d}_2(\mathbf{x})$ *converges to the Bayes risk as* $min\{n_1, \ldots, n_J\} \to \infty$*.*

3.3 Comparison with other depth based classifiers

To compare the performance of our $\mathrm{L}_p\mathrm{D}$ classifiers with existing classifiers based on MD, HD and PD, we carried out some simulations. To keep our examples simple, here we used two-class problems in two dimensions. For varying choices of $(p_1, \mathbf{b}_1, \mathbf{A}_1)$ and $(p_2, \mathbf{b}_2, \mathbf{A}_2)$, we had different types of classification problems (see Table 1). In each case, taking equal number of observations from two competing classes, we generated training and test sets of sizes 400 and 1000, respectively. Each experiment was repeated 200 times, and the average test set misclassification rates of different classifiers were computed over these 200 trials. Note that our proposed method needs $\mathcal{P}$ to be specified. We observed that for higher values of p, l_p contours do not change much (see, e.g., Figure 1, where l_5 contours almost look like l_∞ contours). So, throughout this article, we used $\mathcal{P} = \{2^{(i-1)/2};\ i = 1, 2, \ldots, 10\}$ for our numerical work.

Following Ghosh and Hall (2008), we computed regret functions (i.e., difference between the misclassification rate of a classifier and the Bayes risk) for different classifiers. Table 1 shows the corresponding regret ratios for classifiers based on HD, PD and MD. Regret ratio (η_t) of the classifier t is given by ratio of the regret of that classifier and that of the $\mathrm{L}_p\mathrm{D}$ classifier. Clearly, $\eta_t < 1 (> 1)$ implies that the classifier is better (worse) than the $\mathrm{L}_p\mathrm{D}$ classifier, and the deviation from 1 gives an idea of how better (worse) it is. Since HD and PD classifiers are robust, here we used robust version of MD based on MCD estimates. Following Hubert and van Driessen

Table 1: Regret ratios of other depth-based classifiers.

Choice ↓	Difference in locations				Difference in scales				Difference in shapes			
$(p_1, \mathbf{b}_1, \mathbf{A}_1)$	$(p_2, \mathbf{b}_2, \mathbf{A}_2)$	HD	PD	MD	$(p_2, \mathbf{b}_2, \mathbf{A}_2)$	HD	PD	MD	$(p_2, \mathbf{b}_2, \mathbf{A}_2)$	HD	PD	MD
$(1, \mathbf{0}, \mathbf{I}_2)$	$(1, \mathbf{1}, \mathbf{I}_2)$	53.01	1.54	1.26	$(1, \mathbf{0}, \frac{1}{9}\mathbf{I}_2)$	2.59	1.13	1.01	$(2, \mathbf{0}, \mathbf{I}_2)$	2.52	2.49	1.75
$(2, \mathbf{0}, \mathbf{I}_2)$	$(2, \mathbf{1}, \mathbf{I}_2)$	30.19	2.97	0.96	$(2, \mathbf{0}, \frac{1}{9}\mathbf{I}_2)$	5.88	1.50	0.89	$(4, \mathbf{0}, \mathbf{I}_2)$	1.57	1.40	1.14
$(8, \mathbf{0}, \mathbf{I}_2)$	$(8, \mathbf{1}, \mathbf{I}_2)$	25.94	6.31	1.17	$(8, \mathbf{0}, \frac{1}{9}\mathbf{I}_2)$	5.84	3.93	2.06	$(1, \mathbf{0}, \mathbf{I}_2)$	14.62	2.85	1.67

$\mathbf{0} = (0, 0)^T$, $\mathbf{1} = (1, 1)^T$, $\mathbf{I}_2$ is the 2×2 identity matrix.

(2004), 75% observations were used to compute these estimates. These MCD estimates were also used as $\hat{\boldsymbol{\mu}}$ and $\hat{\boldsymbol{\Sigma}}$ to compute $\hat{\mathbf{A}}$, $\hat{\mathbf{b}}$ and hence the empirical versions of L_p depths. For the location problems, we used maximum depth classifiers based on different notions on depth. In other examples, generalized depth based classifiers were used. In this section, we used equal prior probabilities for two competing classes.

In all these examples, L_pD classifiers outperformed the classifiers based on HD and PD. The classifiers based on HD had the worst performance in all cases, especially in location problems. The discrete nature of the empirical version of HD affected its performance, and it often had issues with ties and zero depth (also see Ghosh and Chaudhuri (2005)). As expected, the classifiers based on MD had a slight edge over the L_pD classifiers when both of the competing classes were l_2-symmetric. But, in all other cases, they were outperformed by the L_pD classifiers.

4 Results from the Analysis of Benchmark Datasets

We analyzed eight benchmark data sets for further assessment of the generalized L_pD classifier. The hemophilia data set was taken from Johnson and Wichern (1992). All other data sets were taken either from the UCI Machine Learning Repository (http://archive.ics.uci.edu/ml/) or from the CMU Datasets Archive (http://lib.stat.cmu.edu/datasets/). Descriptions of these data sets are available at these sources. In the case of blood transfusion data set, following Li et al. (2012), we considered only one of the two linearly dependent variables for our analysis. For the synthetic data and the satellite image (satimage) data, the training and the test sets are well specified. In all other cases, we formed these sets by randomly partitioning the data in a way such that the proportion of different classes in training and test sets were as close as possible. In cases of hemophilia data

and diabetes data, we used training samples of size 50 and 100, respectively. In all other cases, we divided the data set into two nearly equal halves to form the training and the test sets. In each case, this random partitioning was done 500 times. Average misclassification rates of different classifiers were computed over these 500 test sets, and they are reported in Table 3 along with their corresponding standard errors. In cases of synthetic data and satimage data, when a classifier led to a test set misclassification rate Δ, its standard error was computed as $\sqrt{\Delta(1-\Delta)/n_t}$, for n_t being the size of the test set.

To facilitate comparison, along with the misclassification rates of different depth based classifiers, results are also reported for two parametric classifiers (linear discriminant analysis (LDA) and quadratic discriminant analysis (QDA)) and two nonparametric classifiers (kernel discriminant analysis (KDA) and nearest neighbor classifier (k-NN) with the smoothing parameter chosen by the leave-one-out cross-validation method). In some of these data sets, the measurement variables were not of comparable units and scales. So, we used KDA and k-NN (see, e.g., Hastie, Tibshirani and Friedman, 2009; Duda, Hart and Stork, 2012) on the standardized data set, where the moment based estimate of the pooled dispersion matrix was used for standardization. Therefore, to keep our comparisons fair, instead of MCD estimates, here we used moment based estimates of location and scatter parameters for MD and L_pD classifiers as well. Since the classifier based on HD had the highest misclassification rates in almost all cases, we do not report them here. For all these benchmark data sets, sample proportions of different classes were used as their prior probabilities.

Table 2: Average misclassification rates (in %) of different classifiers and their standard errors.

Data	d	J	Train	Test	LDA	QDA	k-NN	KDA	PD	MD	L_pD
Synthetic†	2	2	250	1000	10.80 (.98)	10.20 (.96)	11.70 (1.0)	11.00 (.99)	10.80 (.98)	10.30 (.96)	09.40 (.92)
Hemophilia	2	2	50	25	15.22 (.27)	15.47 (.26)	15.79 (.30)	15.11 (.27)	17.36 (.30)	14.33 (.29)	13.99 (.29)
Blood Tran.	3	2	374	374	22.90 (.03)	22.48 (.05)	21.39 (.06)	22.66 (.05)	22.27 (.08)	22.58 (.06)	22.41 (.07)
Diabetes	5	3	100	45	10.46 (.18)	09.39 (.18)	10.04 (.18)	11.16 (.19)	10.38 (.23)	09.23 (.17)	09.52 (.18)
Pima	8	2	384	384	23.37 (.07)	26.02 (.08)	25.73 (.08)	26.57 (.07)	29.96 (.09)	25.47 (.08)	25.34 (.08)
Vehicle	18	4	423	423	22.49 (.07)	16.38 (.07)	21.84 (.08)	21.45 (.07)	42.64 (.11)	16.29 (.07)	16.16 (.07)
Wisconsin	30	2	284	285	04.71 (.05)	04.74 (.05)	09.36 (.07)	10.15 (.07)	07.45 (.06)	04.92 (.05)	04.63 (.05)
Satimage†	36	6	4435	2000	16.02 (.82)	14.11 (.78)	16.89 (.84)	19.71 (.89)	19.30 (.88)	15.22 (.80)	15.13 (.80)

†Data sets with specific training and test sets.

The overall performance of the generalized L_pD classifier was fairly satisfactory (see Table 2). Except for the blood transfusion data, in all other data sets, it had lower misclassification rates than the PD classifier, and in most of the cases, the difference between their misclassification rates was found to be statistically significant at 5% level when the usual large sample test was used for testing the equality of proportions. In 7 out of 8 data sets, it had lower misclassification rates than the MD classifier as well. It performed better than LDA in all cases barring the Pima Indian data and better than QDA in 6 out of 8 data sets. In cases of haemophilia data and Pima Indian data, its performance was significantly better than QDA. In several data sets, the generalized L_pD and MD classifiers performed better than KDA and *k*-NN classifiers as well. In high-dimensional data sets, especially in vehicle data and Wisconsin breast cancer (diagnostic) data, while nonparmetric classifiers like KDA and *k*-NN had poor performance due to data sparsity, generalized MD and L_pD classifiers were not affected much by curse of dimensionality. In these data sets, these two depth based classifiers yielded substantially lower misclassification rates.

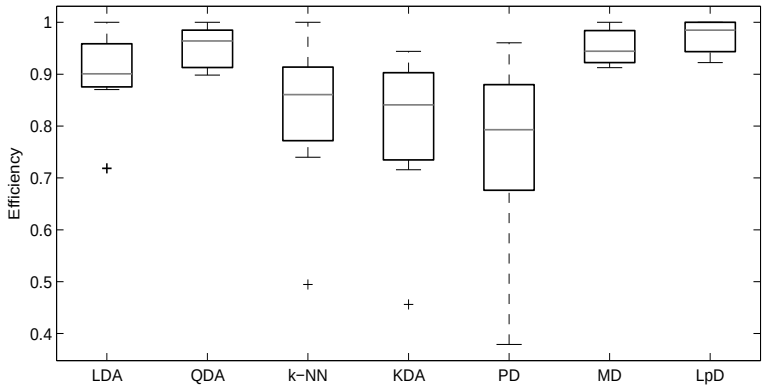

Figure 3: Overall performance of different classifiers on benchmark data sets.

For visual comparison of the overall performance among different classifiers, we computed their efficiencies for different data sets, and they are presented using boxplots in Figure 3. In a particular data set, the efficiency of a classifier t is given by $e_t = \Delta_0/\Delta_t$, where Δ_t is the misclassification rate of the classifier t and $\Delta_0 = \min_t \Delta_t$. Clearly, $e_t = 1$ for the best classifier, and $e_t < 1$ for all other classifiers. Smaller values of e_t indicates the lack of efficiency of the classifier t. Figure 3 clearly shows the superiority of the L_pD classifier over all other classifiers considered here.

Appendix: Proofs and Mathematical Details

Lemma 1: If the density $f(\mathbf{x}, p)$ is of the form $\psi(\|\mathbf{A}(\mathbf{x}-\mathbf{b})\|_p)$ for some $p > 0$ and a continuous function $\psi : \mathbb{R}_+ \to \mathbb{R}_+$, then it can be expressed as

$$f(\mathbf{x}, p) = |\mathbf{A}| \, \frac{p^{d-1}\Gamma(d/p)}{2^d\{\Gamma(1/p)\}^d} \frac{g_p(\delta_p(\mathbf{x}))\delta_p(\mathbf{x})^{d+1}}{[1-\delta_p(\mathbf{x})]^{d-1}} \text{ for } 0 < \delta_p(\mathbf{x}) < 1,$$

where $\delta_p(\mathbf{x}) = [1 + \|\mathbf{A}(\mathbf{x}-\mathbf{b})\|_p]^{-1}$, and g_p is the density function of $\delta_p(\mathbf{X})$ when $\mathbf{X} \sim f(\cdot, p)$.

Proof of Lemma 1: Define $\mathbf{Y} = \mathbf{A}(\mathbf{X}-\mathbf{b})$. If f_0 denotes the density of $\mathbf{Y}$, it is easy to see that $f_0(\mathbf{y}) = |\mathbf{A}|^{-1}\psi(\|\mathbf{y}\|_p)$. Now, following the proof of Lemma 1.4 in Fang, Kotz and Ng (1989), one can show that for any non-negative measurable function ξ, we have

$$\int_{\mathbb{R}^d} \xi(\|\mathbf{y}\|_p)\, d\mathbf{y} = \frac{2^d\{\Gamma(1/p)\}^d}{p^d\Gamma(d/p)} \int_0^\infty \xi(u^{1/p}) u^{d/p-1} du.$$

So, for any non-negative measurable function ϕ, defining $\xi = \phi.\psi$ we get

$$\begin{aligned} E[\phi(\|\mathbf{Y}\|_p)] &= |\mathbf{A}|^{-1} \int \phi(\|\mathbf{y}\|_p)\psi(\|\mathbf{y}\|_p)\, d\mathbf{y} = |\mathbf{A}|^{-1} \frac{2^d\{\Gamma(1/p)\}^d}{p^d\Gamma(d/p)} \int_0^\infty \phi(u^{1/p})\, \psi(u^{1/p})\, u^{d/p-1}\, du \\ &= |\mathbf{A}|^{-1} \frac{2^d\{\Gamma(1/p)\}^d}{p^{d-1}\Gamma(d/p)} \int_0^\infty \phi(r)\, \psi(r)\, r^{d-1}\, dr \text{ (letting } r = u^{1/p}). \end{aligned}$$

Therefore, the density of $r_p(\mathbf{X}) = \|\mathbf{A}(\mathbf{X}-\mathbf{b})\|_p$ is the form $\tau_p(r) = |\mathbf{A}|^{-1} \frac{2^d\{\Gamma(1/p)\}^d}{p^{d-1}\Gamma(d/p)}\, \psi(r)\, r^{d-1}$, which implies $f(\mathbf{x}, p) = \psi(r_p(\mathbf{x})) = |\mathbf{A}| \, \frac{p^{d-1}\Gamma(d/p)}{2^d\{\Gamma(1/p)\}^d} \, \frac{\tau_p(r_p(\mathbf{x}))}{r_p(\mathbf{x})^{d-1}}$. Recall that $\delta_p(\mathbf{x}) = [1 + r_p(\mathbf{x})]^{-1}$, and using the change of variables formula, we have the final expression for the density $f(\mathbf{x}, p)$ in terms of $\delta_p(\mathbf{x})$ and its density $g_p(\cdot)$. □

Corollary 1: Define $\tilde{\delta}_p(\mathbf{x}) = [1 + a_0 r_p(\mathbf{x})]^{-1}$, where a_0 is a positive constant, and let $g_p^{(a)}(\cdot)$ denote the density of $\tilde{\delta}_p(\mathbf{X})$. If the density $f(\mathbf{x}, p)$ is of the form $\psi(\|\mathbf{A}(\mathbf{x}-\mathbf{b})\|_p)$ for some $p > 0$ and a continuous function $\psi : \mathbb{R}_+ \to \mathbb{R}_+$, it can be expressed as

$$f(\mathbf{x}, p) = a_0^d \, |\mathbf{A}| \, \frac{p^{d-1}\Gamma(d/p)}{2^d\{\Gamma(1/p)\}^d} \frac{g_p^{(a)}(\tilde{\delta}_p(\mathbf{x}))\tilde{\delta}_p(\mathbf{x})^{d+1}}{[1-\tilde{\delta}_p(\mathbf{x})]^{d-1}} \text{ for } 0 < \tilde{\delta}_p(\mathbf{x}) < 1,$$

Proof of Corollary 1: If we define $\tilde{r}_p(\mathbf{x}) = a_0 r_p(\mathbf{x})$ and if $\tau_p^{(a)}(\cdot)$ denotes the density of $\tilde{r}_p(\mathbf{X})$, it

is easy to check that $\frac{\tau_p(r_p(\mathbf{x}))}{r_p(\mathbf{x})^{d-1}} = a_0^d \frac{\tau_p^{(a)}(\tilde{r}_p(\mathbf{x}))}{\tilde{r}_p(\mathbf{x})^{d-1}}$. Now from the proof of Lemma 1, it follows that $f(\mathbf{x}, p) = a_0^d\,|\mathbf{A}|\,\frac{p^{d-1}\Gamma(d/p)}{2^d\{\Gamma(1/p)\}^d}\,\frac{\tau_p^{(a)}(\tilde{r}_p(\mathbf{x}))}{\tilde{r}_p(\mathbf{x})^{d-1}}$. Since $\tilde{\delta}_p(\mathbf{x}) = [1+\tilde{r}_p(\mathbf{x})]^{-1}$, the result is obtained by using the change of variables formula. □

Lemma 2: Let f be l_{p_0}-symmetric for some $p_0 \geq 1$. Define $\tilde{\delta}_p(\mathbf{x})$ and $g_p^{(a)}$ as in Corollary 1.

(i) If $\hat{\mathbf{b}}$ and $\hat{\mathbf{A}}$ satisfy assumption $(C1)$, for any $p \geq 1$, $\sup_{\mathbf{x}\in\mathbb{R}^d} |\hat{\delta}_p(\mathbf{x}) - \tilde{\delta}_p(\mathbf{x})| = O_P(n^{-1/2})$.

(ii) Further, if the density g_p and the bandwidth h associated with the kernel density estimator $\hat{g}_{p,h}(\cdot)$ satisfies assumption $(C2)$, then $\sup_{\mathbf{x}\in\mathbb{R}^d} |\hat{g}_{p,h}(\hat{\delta}_p(\mathbf{x})) - g_p^{(a)}(\tilde{\delta}_p(\mathbf{x}))| \xrightarrow{P} 0$ as $n \to \infty$.

Proof of Lemma 2: (i) For any matrix $\mathbf{A}$, let $||\mathbf{A}||_F = [trace(\mathbf{A}^T\mathbf{A})]^{1/2}$ denote its Frobenius norm and $||\mathbf{A}||_p = \sup_{\mathbf{x}}\{||A\mathbf{x}||_p/||\mathbf{x}||_p\}$ denote its p-th norm. For any $p \geq 1$, $||\mathbf{A}||_p$ is an induced norm that satisfies $||\mathbf{A}\mathbf{x}||_p \leq ||\mathbf{A}||_p||\mathbf{x}||_p$ for all $\mathbf{x} \in \mathbb{R}^d$. In finite dimension, since all norms are equivalent, we have $||\mathbf{A}\mathbf{x}||_p \leq C_p||\mathbf{A}||_F||\mathbf{x}||_p$ for some positive constant C_p. This also implies that $||\mathbf{x}||_p \leq C_p||\mathbf{A}^{-1}||_F||\mathbf{A}\mathbf{x}||_p$, or $||\mathbf{A}\mathbf{x}||_p \geq ||\mathbf{x}||_p/C_p||\mathbf{A}^{-1}||_F$.

Let us divide $\mathbb{R}^d$ into two disjoint regions: $\mathcal{R}_1 = \{\mathbf{x} : ||\mathbf{x}|| \leq M\}$ and $\mathcal{R}_2 = \{\mathbf{x} : ||\mathbf{x}|| > M\}$, where M is a large positive constant. Note that $|\hat{\delta}_p(\mathbf{x}) - \tilde{\delta}_p(\mathbf{x})| = \frac{|\hat{r}_p(\mathbf{x}) - a_0 r_p(\mathbf{x})|}{[1+\hat{r}_p(\mathbf{x})][1+a_0 r_p(\mathbf{x})]} \leq |\hat{r}_p(\mathbf{x}) - a_0 r_p(\mathbf{x})|$. Using the triangle inequality, we obtain

$$\sqrt{n}\sup_{\mathbf{x}\in\mathcal{R}_1}|\hat{\delta}_p(\mathbf{x}) - \tilde{\delta}_p(\mathbf{x})| \leq \sqrt{n}\sup_{\mathbf{x}\in\mathcal{R}_1}|\hat{r}_p(\mathbf{x}) - a_0 r_p(\mathbf{x})| \leq ||\hat{\mathbf{A}}||_F\sqrt{n}||\hat{\mathbf{b}} - \mathbf{b}||_p + \sqrt{n}||\hat{\mathbf{A}} - a_0\mathbf{A}||_F \sup_{\mathbf{x}\in\mathcal{R}_1}||\mathbf{x} - \mathbf{b}||_p.$$

Since $\sup_{\mathbf{x}\in\mathcal{R}_1}||\mathbf{x} - \mathbf{b}||_p \leq M + ||\mathbf{b}||$, under the condition $(C1)$, $\sup_{\mathbf{x}\in\mathcal{R}_1}\sqrt{n}|\hat{\delta}_p(\mathbf{x}) - \tilde{\delta}_p(\mathbf{x})| = O_p(1)$. Also note that $|\hat{\delta}_p(\mathbf{x}) - \tilde{\delta}_p(\mathbf{x})| \leq \frac{|\hat{r}_p(\mathbf{x}) - a_0 r_p(\mathbf{x})|}{[\hat{r}_p(\mathbf{x})][a_0 r_p(\mathbf{x})]}$ and this implies that

$$\sqrt{n}|\hat{\delta}_p(\mathbf{x}) - \tilde{\delta}_p(\mathbf{x})| \leq \frac{||\hat{\mathbf{A}}||_F\sqrt{n}||\hat{\mathbf{b}} - \mathbf{b}||_p}{(||\hat{\mathbf{A}}\mathbf{x}||_p - ||\hat{\mathbf{A}}\hat{\mathbf{b}}||_p)(||a_0\mathbf{A}\mathbf{x}||_p - ||a_0\mathbf{A}\mathbf{b}||_p)} + \frac{\sqrt{n}||\hat{\mathbf{A}} - (a_0\mathbf{A})||_F}{(||\hat{\mathbf{A}}\mathbf{x}||_p - ||\hat{\mathbf{A}}\hat{\mathbf{b}}||_p)a_0||\mathbf{A}||_F}.$$

Therefore, $\sqrt{n}\sup_{\mathbf{x}\in\mathcal{R}_2}|\hat{\delta}_p(\mathbf{x}) - \tilde{\delta}_p(\mathbf{x})|$

$$\leq \frac{||\hat{\mathbf{A}}||_F\sqrt{n}||\hat{\mathbf{b}} - \mathbf{b}||_p}{[(M/||\hat{\mathbf{A}}^{-1}||_F) - ||\hat{\mathbf{A}}\hat{\mathbf{b}}||_p][(M/||C_p a_0\mathbf{A}^{-1}||_F) - ||a_0\mathbf{A}\mathbf{b}||_p]} + \frac{\sqrt{n}||\hat{\mathbf{A}} - (a_0\mathbf{A})||_F}{[(M/||\hat{\mathbf{A}}^{-1}||_F) - ||\hat{\mathbf{A}}\hat{\mathbf{b}}||_p]a_0||\mathbf{A}||_F}.$$

In each of the two terms on the right side, the numerator is $O_p(1)$, while the denominator coverges to a positive constant. So, we have $\sqrt{n}\sup_{\mathbf{x}\in\mathcal{R}_2}|\hat{\delta}_p(\mathbf{x}) - \tilde{\delta}_p(\mathbf{x})| = O_p(1)$. Now, combining the results

on $\mathcal{R}_1$ and $\mathcal{R}_2$, the first part of the lemma is proved.

(ii) Note that $|\hat{g}_{p,h}(\hat{\delta}_p(\mathbf{x})) - g_p^{(a)}(\tilde{\delta}_p(\mathbf{x}))| \leq |\hat{g}_{p,h}(\hat{\delta}_p(\mathbf{x})) - \hat{g}_{p,h}(\tilde{\delta}_p(\mathbf{x}))| + |\hat{g}_{p,h}(\tilde{\delta}_p(\mathbf{x})) - g_p^{(a)}(\tilde{\delta}_p(\mathbf{x}))|$. An application of the triangle inequality implies that

$$|\hat{g}_{p,h}(\hat{\delta}_p(\mathbf{x})) - \hat{g}_{p,h}(\tilde{\delta}_p(\mathbf{x}))| \leq \frac{1}{nh}\sum_{i=1}^{n}\left|K\left[\frac{\hat{\delta}_p(\mathbf{x}) - \hat{\delta}_p(\mathbf{x}_i)}{h}\right] - K\left[\frac{\tilde{\delta}_p(\mathbf{x}) - \hat{\delta}_p(\mathbf{x}_i)}{h}\right]\right|,$$

where K is the kernel function associated with the density estimate. Using the mean value theorem and assuming $C_K = \sup_x |K'(x)| < \infty$ (here we use the Gaussian kernel, which has bounded first derivative), we have $\sup_{\mathbf{x}\in\mathbb{R}^d} |\hat{g}_{p,h}(\hat{\delta}_p(\mathbf{x})) - \hat{g}_{p,h}(\tilde{\delta}_p(\mathbf{x}))| \leq C_K \sup_{\mathbf{x}\in\mathbb{R}^d} |\hat{\delta}_p(\mathbf{x}) - \tilde{\delta}_p(\mathbf{x})|/h^2$. So, if h satisfies the condition $(C2)$, $\sup_{\mathbf{x}\in\mathbb{R}^d} |\hat{g}_{p,h}(\hat{\delta}_p(\mathbf{x})) - \hat{g}_{p,h}(\tilde{\delta}_p(\mathbf{x}))| \xrightarrow{P} 0$ as $n \to \infty$.

To prove the convergence of $\sup_{\mathbf{x}\in\mathbb{R}^d} |\hat{g}_{p,h}(\tilde{\delta}_p(\mathbf{x})) - g_p^{(a)}(\tilde{\delta}_p(\mathbf{x}))|$, first note that

$$\sup_{\mathbf{x}\in\mathbb{R}^d} |\hat{g}_{p,h}(\tilde{\delta}_p(\mathbf{x})) - g_p^{(a)}(\tilde{\delta}_p(\mathbf{x}))| \leq \sup_{\mathbf{x}\in\mathbb{R}^d} |\hat{g}_{p,h}(\tilde{\delta}_p(\mathbf{x})) - g_{p,h}^{*}(\tilde{\delta}_p(\mathbf{x}))| + \sup_{\mathbf{x}\in\mathbb{R}^d} |g_{p,h}^{*}(\tilde{\delta}_p(\mathbf{x})) - g_p^{(a)}(\tilde{\delta}_p(\mathbf{x}))|,$$

where $g_{p,h}^{*}(\tilde{\delta}_p(\mathbf{x})) = (nh)^{-1}\sum_{i=1}^{n} K[\{\tilde{\delta}_p(\mathbf{x}) - \tilde{\delta}_p(\mathbf{x}_i)\}/h]$. So, using triangle inequality, one gets

$$\sup_{\mathbf{x}\in\mathbb{R}^d} |\hat{g}_{p,h}(\tilde{\delta}_p(\mathbf{x})) - g_p^{(a)}(\tilde{\delta}_p(\mathbf{x})) \leq \sup_{\mathbf{x}\in\mathbb{R}^d} |g_{p,h}^{*}(\tilde{\delta}_p(\mathbf{x})) - g_p^{(a)}(\tilde{\delta}_p(\mathbf{x}))| + C_K \sup_{\mathbf{x}\in\mathbb{R}^d} |\hat{\delta}_p(\mathbf{x}) - \tilde{\delta}_p(\mathbf{x})|/h^2.$$

We have already proved the convergence of the second term on the right side to 0. Since g_p is uniformly continuous and $nh/\log(n) \to \infty$ as $n \to \infty$ (see assumption (C2)), the convergence of the first term and hence the result now follow from the uniform convergence property of the kernel density estimate $g_{p,h}^{*}$ (see, e.g., Silverman, 1998).

Proof of Theorem 1: Define $A_n = \{\mathbf{x} : \tilde{\delta}_p(\mathbf{x}) \in [\zeta_{1n}, \zeta_{2n}]\} \subset \mathbb{R}^d$. Using the mean value theorem on the logarithmic function, one gets

$$\sup_{\mathbf{x}\in A_n} |\ln \hat{g}_{p,h}(\hat{\delta}_p(\mathbf{x})) - \ln g_p^{(a)}(\tilde{\delta}_p(\mathbf{x}))| \leq \sup_{\mathbf{x}\in\mathbb{R}^d} |\hat{g}_{p,h}(\hat{\delta}_p(\mathbf{x})) - g_p^{(a)}(\tilde{\delta}_p(\mathbf{x}))| \sup_{\mathbf{x}\in A_n}\left\{\frac{1}{\xi_n(\mathbf{x})}\right\},$$

where $\xi_n(\mathbf{x}) = \lambda_n \hat{g}_{p,h}(\hat{\delta}_p(\mathbf{x})) + (1-\lambda_n) g_p^{(a)}(\tilde{\delta}_p(\mathbf{x}))$ for some $\lambda_n \in (0,1)$. Hence, $\xi_n(\mathbf{x})$ converges to $g_p^{(a)}(\tilde{\delta}_p(\mathbf{x}))$ as $n \to \infty$ (follows from Lemma 2). From the proof of Lemma 2 and using a result on the rate of uniform consistency of the kernel density estimate (see Theorem B in Silverman, 1978, p. 181). Under (C2), we have $\sup_{\mathbf{x}\in\mathbb{R}^d} |\hat{g}_{p,h}(\hat{\delta}_p(\mathbf{x})) - g_p^{(a)}(\tilde{\delta}_p(\mathbf{x}))| = O_P(n^{-1/2}h^{-2})$. Therefore, using

$(C3)$, we get $\sup_{\mathbf{X}\in A_n} |\ln \hat{g}_{p,h}(\hat{\delta}_p(\mathbf{x})) - \ln g_p^{(a)}(\tilde{\delta}_p(\mathbf{x}))| \xrightarrow{P} 0$ as $n \to \infty$.

Following similar arguments as above, one can show that

$$\sup_{\mathbf{x}\in A_n} |\ln \hat{\delta}_p(\mathbf{x}) - \ln(\tilde{\delta}_p(\mathbf{x}))| \leq \sup_{\mathbf{x}\in\mathbb{R}^d} |\hat{\delta}_p(\mathbf{x}) - (\tilde{\delta}_p(\mathbf{x}))| \sup_{\mathbf{x}\in A_n} \left\{\frac{1}{\gamma_n(\mathbf{x})}\right\},$$

where $\gamma_n(\mathbf{x}) = \lambda_n\hat{\delta}_p(\mathbf{x}) + (1-\lambda_n)\tilde{\delta}_p(\mathbf{x})$ for some $\lambda_n \in (0,1)$, and it converges to $\tilde{\delta}_p(\mathbf{x})$. Since $\sup_{\mathbf{x}\in\mathbb{R}^d} |\hat{\delta}_p(\mathbf{x}) - (\tilde{\delta}_p(\mathbf{x}))| = O_P(n^{-1/2})$ and $n^{1/2}\zeta_{1n} \to \infty$ as $n \to \infty$, we have $\sup_{\mathbf{X}\in A_n} |\log \hat{\delta}_p(\mathbf{x}) - \log(\tilde{\delta}_p(\mathbf{x}))| \xrightarrow{P} 0$ as $n \to \infty$. Similarly, using the fact that $n^{1/2}(1-\zeta_{2n}) \to \infty$ as $n \to \infty$, under the given conditions, one can show that $\sup_{\mathbf{X}\in A_n} |\log(1-\hat{\delta}_p(\mathbf{x})) - \log(1-\tilde{\delta}_p(\mathbf{x}))| \xrightarrow{P} 0$ as $n \to \infty$. Combining these results and using corollary 1, we get $\sup_{\mathbf{X}\in A_n} |\ln \hat{f}_n(\mathbf{x}, p) - \ln(a_0^{-d} f(\mathbf{x}, p))| \xrightarrow{P} 0$ as $n \to \infty$.

Define the set $I_n = \{i : \mathbf{X}_i \in A_n\}$ and note that $\left|\frac{1}{n}\sum_{i\in I_n} \ln \hat{f}_n(\mathbf{x}_i, p) - E_{p_0}\{\ln(a_0^{-d} f(\mathbf{X}, p))\}\right| \leq \sup_{\mathbf{X}\in A_n} |\ln \hat{f}_n(\mathbf{x}, p) - \ln(a_0^{-d} f(\mathbf{x}, p))| + \left|\frac{1}{n}\sum_{i\in I_n} \ln(a_0^{-d} f(\mathbf{x}_i, p)) - E_{p_0}\{\ln(a_0^{-d} f(\mathbf{X}, p))\}\right|$, where E_{p_0} denotes the expectation with respect to $f(\mathbf{x}, p_0)$. We have already proved the probability convergence of the first part on the right side to 0. From the strong law of large numbers, we have $\left|\frac{1}{n}\sum_{i\in I_n} \ln(a_0^{-d} f(\mathbf{x}_i, p)) - E_{p_0}\{\ln(a_0^{-d} f(\mathbf{X}, p))I(\mathbf{X}\in A_n)\}\right| \xrightarrow{a.s.} 0$ as $n \to \infty$.

Using monotone convergence theorem on the positive and negative parts of the integrand separately, we have $E_{p_0}\{\ln(a_0^{-d} f(\mathbf{X}, p))I(\mathbf{X}\in A_n)\} \to E_{p_0}\{\ln(a_0^{-d} f(\mathbf{X}, p))\}$ as $n \to \infty$. So, for all $p \geq 1$, we get $\frac{1}{n}\sum_{i\in I_n} \ln \hat{f}_n(\mathbf{x}_i, p) \xrightarrow{P} E_{p_0}\{\ln(a_0^{-d} f(\mathbf{X}, p))\} = \ln(a_0^{-d}) + E_{p_0}\{\ln f(\mathbf{X}, p)\}$. One can notice that $E_{p_0}\{\ln f(\mathbf{X}, p)\} = E_{p_0}\{\ln f(\mathbf{X}, p_0)\} - KL(p_0, p)$, where $KL(p_0, p)$ denotes the Kullback-Leibler divergence between $f(\mathbf{X}, p_0)$ and $f(\mathbf{X}, p)$. So, maximization of $E_{p_0}\{\ln f(\mathbf{X}, p)\}$ is equivalent to minimization of $KL(p_0, p)$ over $\mathcal{P}$. Using Jensen's inequality on the logarithmic function, we also have $E_{p_0}\{\ln f(\mathbf{X}, p)\} < E_{p_0}\{\ln f(\mathbf{X}, p_0)\}$, or $KL(p_0, p) > KL(p_0, p_0)$ for all $p \neq p_0$. Since $\mathcal{P}$ is a finite set, the proof now follows from the uniqueness of the minimizer of $KL(p_0, p)$. □

Proof of Theorem 2: The misclassification rate of the classifier d_1 is given by

$$\Delta(\mathsf{d}_1) = \frac{1}{J}\sum_{j=1}^{J} \int \theta_j(\mathbf{x}, \hat{p}_n) f_j(\mathbf{x}, p_0)\, d\mathbf{x},$$

where $\theta_j(\mathbf{x}, \hat{p}_n) = P\{\hat{\delta}_{\hat{p}_n,j}(\mathbf{x}) < \hat{\delta}_{\hat{p}_n,i}(\mathbf{x}) \text{ for some } i \neq j\}$. From the arguments given in the proof of

Theorem 1, it follows that $\hat{p}_n \xrightarrow{P} p_0$. Using the continuous mapping theorem, one gets $\hat{\delta}_{\hat{p}_n,j}(\mathbf{x}) \xrightarrow{P} \tilde{\delta}_{p_0,j}(\mathbf{x})$ for all $1 \le j \le J$. So, if $\tilde{\delta}_{p_0,j}(\mathbf{x}) < \tilde{\delta}_{p_0,i}(\mathbf{x})$ for all $i \neq j$ (or equivalently, $f_j(\mathbf{x}, p_0) > f_i(\mathbf{x}, p_0)$ for all $i \neq j$), $\theta_j(\mathbf{x}, \hat{p}_n)$ converges to 0, otherwise it converges to 1. Therefore, using the Dominated Convergence Theorem, we have the convergence of $\Delta(\mathsf{d}_1)$ to the Bayes risk. □

Lemma 3: If $p_0 \in \mathcal{P}$, under $(C1)$ - $(C3)$, for any fixed $\mathbf{x} \in \mathbb{R}^d$, $\hat{f}_h(\mathbf{x}, p_0) \xrightarrow{P} a_0^{-d} f(\mathbf{x}, p_0)$ as $n \to \infty$.

Proof of Lemma 3: Recall that $\hat{f}_h(\mathbf{x}, p_0) = C_{\hat{p}_0,d} \dfrac{\hat{g}_{p_0,h}(\hat{\delta}_{p_0}(\mathbf{x}))(\hat{\delta}_{p_0}(\mathbf{x}))^{d+1}}{(1-\hat{\delta}_{p_0}(\mathbf{x}))^{d-1}}$, where $C_{p,d} = \dfrac{p^{d-1}\Gamma(d/p)}{2^d\{\Gamma(1/p)\}^d}$. From the continuous mapping theorem and Theorem 1, we get $(\hat{\delta}_{\hat{p}_n}(\mathbf{x}))^{d+1} \xrightarrow{P} (\tilde{\delta}_{p_0}(\mathbf{x}))^{d+1}$, $(1 - \hat{\delta}_{\hat{p}_n}(\mathbf{x}))^{d-1} \xrightarrow{P} (1-\tilde{\delta}_{p_0}(\mathbf{x}))^{d-1}$ for $d > 1$ and $\hat{g}_{\hat{p}_0,h}(\delta_{p_0}(\mathbf{x})) \xrightarrow{p} g_{p_0}^{(a)}(\delta_{p_0}(\mathbf{x}))$. The result now follows from Slutsky's lemma and Corollary 1. □

Proof of Theorem 3 : For simplicity, we consider the case when $J = 2$. For $J > 2$, we use pairwise classification, and hence the result can be obtained by repeating the same argument for each of the $J(J-1)/2$ pairs of classes (see Tewari and Bartlett, 2007). Using Lemma 3, we get $(\hat{f}_{h_1}(\mathbf{x}, p_1^\circ), \hat{f}_{h_2}(\mathbf{x}, p_2^\circ)) \xrightarrow{P} (a_1^{-d} f(\mathbf{x}, p_1^\circ), a_2^{-d} f(\mathbf{x}, p_2^\circ))$, where $a_1, a_2 > 0$. Consider the classifier $\mathsf{d}_2^{p_1^\circ, p_2^\circ}$ which is of the form:

$$\mathsf{d}_2^{p_1^\circ,p_2^\circ}(\mathbf{x}) = \begin{cases} 1 & \text{if } \hat{f}_{h_1}(\mathbf{x}, p_1^\circ)/\hat{f}_{h_2}(\mathbf{x}, p_2^\circ) > k, \\ 2 & \text{otherwise.} \end{cases}$$

Here, k is chosen by minimizing the leave-one-out cross-validation estimate of the misclassification probability. Note that for any fixed k, this cross-validation estimate is given by

$$\Delta^{CV}_{n,p_1^\circ,p_2^\circ}(k) = \sum_{i=1, j\neq i}^{2} \frac{\pi_i}{n_i} \sum_{l=1}^{n_i} I\left[\left\{\hat{f}_{h_i}^{-il}(\mathbf{x}_{il}, p_i^\circ)/\hat{f}_{h_j}^{-il}(\mathbf{x}_{jl}, p_j^\circ)\right\} \le k_i\right],$$

where $n = (n_1, n_2)$, $k_1 = k$, $k_2 = 1/k$, and $\hat{f}_{h_j}^{-il}$ denotes the leave-one-out density estimate with $\mathbf{x}_{il}$ being left out as a training data point. Let k_n denote the minimizer of $\Delta^{CV}_{n,p_1^\circ,p_2^\circ}(k)$.

Define

$$\Delta_{p_1^\circ,p_2^\circ}(k) = \sum_{i=1, j\neq i}^{2} \pi_i P\left[\left\{a_i^{-d} f(\mathbf{X}, p_i^\circ)/a_j^{-d} f(\mathbf{X}, p_j^\circ)\right\} \le k_i \;\middle|\; \mathbf{X} \in i\text{-th class}\right],$$

and $k_0 = \arg\min_k \Delta_{p_1^\circ,p_2^\circ}(k)$. Under the stated assumptions, $k_0 = a_1^{-d}\pi_2/a_2^{-d}\pi_1$ is a unique solution, and $0 < k_0 < \infty$. Note that $\Delta_{p_1^\circ,p_2^\circ}(k_0)$ denotes the misclassification probability of the Bayes clas-

sifier d_B. Consistency of the density estimates and using the same arguments as in Lemma 3 of Dutta and Ghosh (2012), one can show that $\sup_k |\Delta^{CV}_{n,p_1^\circ,p_2^\circ}(k) - \Delta_{p_1^\circ,p_2^\circ}(k)| \xrightarrow{P} 0$. This now implies that $k_n \xrightarrow{P} k_0$ as $\min\{n_1, n_2\} \to \infty$. So, using the continuous mapping theorem and continuity of the underlying densities, we have $\mathsf{d}_2^{p_1^\circ,p_2^\circ}(\mathbf{x}) \xrightarrow{P} \mathsf{d}_B(\mathbf{x})$ for any fixed $\mathbf{x} \in \mathbb{R}^d$ not lying on the Bayes class boundary. Using the Dominated Convergence Theorem, we can show that the misclassification probability of the classifier $\mathsf{d}_2^{p_1^\circ,p_2^\circ}$ has the following convergence:

$$\Delta_2^{p_1^\circ,p_2^\circ} = \sum_{i=1}^{2} \pi_i \, P[\mathsf{d}_2^{p_1^\circ,p_2^\circ}(\mathbf{X}) \neq i \mid \mathbf{X} \sim f(\cdot, p_i^\circ)] \to \sum_{i=1}^{2} \pi_i \, P[\mathsf{d}_B(\mathbf{X}) \neq i \mid \mathbf{X} \sim f(\cdot, p_i^\circ)] = \Delta_B,$$

the Bayes risk as $\min\{n_1, n_2\} \to \infty$.

From Theorem 1, we have $\hat{p}_{1n_1} \xrightarrow{P} p_1^\circ$ and $\hat{p}_{2n_2} \xrightarrow{P} p_2^\circ$ as $n \to \infty$. Since $\mathcal{P}$ is finite, this implies $P(\hat{p}_{1n_1} = p_1^\circ, \hat{p}_{2n_2} = p_2^\circ) \to 1$ as $n \to \infty$. The misclassification probability of the classifier d_2 can be expressed as follows:

$$\sum_{p_1 \in \mathcal{P}} \sum_{p_2 \in \mathcal{P}} \Delta_2^{p_1,p_2} \, P(\hat{p}_{1n_1} = p_1, \hat{p}_{1n_1} = p_2).$$

It is now straight forward to prove that the expression above converges to $\Delta_2^{p_1^\circ,p_2^\circ} (= \Delta_B,$ the Bayes risk) as $\min\{n_1, n_2\} \to \infty$. □

References

[1] Arellano-Valle, R. B. and Richter, W.-D. (2012) On skewed continuous $l_{n,p}$-symmetric distributions. *Chilean J. Statist.*, **3**, 193-212.

[2] Cator, E. A. and Lopuhaä, H. P. (2012) Central limit theorem and influence function for the MCD estimators at general multivariate distributions. *Bernoulli*, **18**, 520-551.

[3] Chakraboty, B. and Chaudhuri, P. (1996) On a transformation and re-transformation technique for constructing affine equivariant multivariate median. *Proc. Amer. Math. Soc.*, **124**, 2539-2547.

[4] Chakraborty, B. (2001) On affine equivariant multivariate quantiles. *Ann. Inst. Statist. Math.*, **53**, 380-403.

[5] Cui, X., Lin, L. and Yang, G. (2008) An extended projection data depth and its applications to discrimination. *Comm. Statist. Theory Methods*, **37**, 2276-2290.

[6] Duda, R. O., Hart, P. E. and Stork, D. G. (2012) *Pattern Classification*, Wiley, New York.

[7] Dutta, S., Ghosh, A. K. and Chaudhuri, P. (2011) Some intriguing properties of Tukey's half-space depth. *Bernoulli*, **17**, 1420-1434.

[8] Dutta, S. and Ghosh, A. K. (2012) On robust classification using projection depth. *Ann. Inst. Statist. Math.*, **64**, 657-676.

[9] Fang, K. -T., Kotz, S. and Ng, K. -W. (1989) *Symmetric Multivariate and Related Distributions*. Chapman and Hall, London.

[10] Fraiman, R., Liu, R. Y. and Meloche, J. (1997) Multivariate density estimation by probing depth. In *L_1-Statistical Procedures and Related Topics. IMS Lecture Notes* (Ed. Y. Dodge), Inst. Math. Statist., Hayward, California, **31**, 415-430.

[11] Ghosh, A. K. and Chaudhuri, P. (2005) On maximum depth and related classifiers. *Scand. J. Statist.*, **32**, 328-350.

[12] Ghosh, A. K. and Hall, P. (2008) On error rate estimation in nonparametric classification. *Statistica Sinica*, **18**, 1081-1100.

[13] Gupta, A. K. and Song, D. (1997) L_p norm spherical distribution. *J. Statist. Plan. Infer.*, **60**, 241-260.

[14] Hartikainen, A. and Oja, H. (2006) On nonparametric discrimination rules. In *Data Depth: Robust Multivariate Analysis, Computational Geometry, and Applications* (Eds. R. Liu, R. Serfling and D. Souvaine), Amer. Math. Soc., **72**, 61-70.

[15] Hastie, T., Tibshirani, R. and Friedman, J. (2009) *Elements of Statistical Learning Theory*. Wiley, New York.

[16] Hubert, M. and van Driessen, K. (2004) Fast and robust discriminant analysis. *Comput. Statist. Data Anal.*, **45**, 301-320.

[17] Hoberg, R. and Mosler, K. (2006) Data analysis and classification with the zonoid depth. In *Data Depth: Robust Multivariate Analysis, Computational Geometry and Applications* (Eds. R. Liu, R. Serfling and D. Souvaine), Amer. Math. Soc., **72**, 49-59.

[18] Johnson, R. A. and Wichern, D. W. (1992) *Applied Multivariate Statistical Analysis*. Prentice-Hall, New Jersey.

[19] Jörnsten, R. (2004) Clustering and classification based on the L_1 data depth. *J. Multivar. Anal.*, **90**, 67-89.

[20] Li, J., Cuesta-Albertos, J. A. and Liu, R. (2012) Nonparametric classification procedures based on DD-plot. *J. Amer. Statist. Assoc.*, **107**, 737-753.

[21] Liu, R., Parelius, J. and Singh, K. (1999) Multivariate analysis of the data-depth : Descriptive statistics and inference. *Ann. Statist.*, **27**, 783-858.

[22] Paindaveine, D. and Van Bever, G. (2015) Nonparametrically consistent depth-based classifiers. *Bernoulli*, **21**, 62-82.

[23] Sheather, S. J. and Jones, M. C. (1991) A reliable data based bandwidth selection method for kernel density estimation. *J. Royal Statist. Soc. Ser. B*, **53**, 683-690.

[24] Silverman, B. W. (1978) Weak and strong uniform consistency of the kernel estimate of a density and its derivatives. *Ann. Statist.*, **6**, 177-184.

[25] Silverman, B. W. (1998) *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, London.

[26] Sinz, F. and Bethge, M. (2010) L_p-nested symmetric distributions. *J. Mach. Learn. Res.*, **11**, 3409-3451.

[27] Tewari, A. and Bartlett, P. L. (2007) On the consistency of multiclass classification methods. *J. Mach. Learn. Res.*, **8**, 1007-1025.

[28] Tukey, J. W. (1975) Mathematics and the picturing of data. *Proc. 1975 Inter. Cong. Math.*, Vancouver, 523-531.

[29] Yue, X. and Ma, C. (1995) Multivariate L_p-norm symmetric distributions. *Statist. Probab. Letters*, **24**, 281-288.

[30] Zuo, Y. and Serfling, R. (2000) General notions of statistical depth function. *Ann. Statist.*, **28**, 461-482.